A low-power digital camera-on-a-chip implemented in CMOS active pixel approach

Bedabrata Pain, Timothy Shaw, Monico Ortiz, Brita Olson, Julie Heynssens, Chris Wrigley, Charlie Ho, and Guang Yang

California Institute of Technology
Jet Propulsion Laboratory
4800 Oak Grove Drive
MS 300-315
Pasadena, CA 91109
Voice: (818) 354-8765
Fax: (818) 393-0045
Email: bpain@jpl.nasa.gov

INTRODUCTION

For more than three decades now, CMOS technology has been the backbone behind the phenomenal advances in low-cost, low-power, reliable, highly integrated and miniaturized system development. These systems have found widespread use in consumer electronics, micro-processors, memory, audio systems, and radio-frequency circuits (e.g. cellular phones). However, video technology has remained outside the CMOS technology mainstream. Over the years, the incumbent technology in electronic imaging for consumer, commercial, and scientific applications has been charge coupled devices (CCD). The predominance of CCD in electronic imaging was due to its superior sensitivity, dynamic range, uniformity, low noise, and small pixel size. However in order to achieve high (0.999999) charge transfer efficiency, CCDs require specialized silicon processing that is not compatible with CMOS technology. Furthermore, CCDs are high capacitance devices, requiring multiple non-standard and high voltage clocks and biases, while providing only serial output. High device capacitance, large clock swing, need for DC-DC converters, and inability to integrate control & processing electronics on the imager chip make the CCD-based imaging system bulky and power-hungry (camcorder CCD power dissipation is around 10 W). Incompatibility with CMOS technology is a major barrier to realizing low-power, low-cost, digital, integrated system-on-a-chip using CCDs.

Despite several efforts in the decade of eighties and earlier, CMOS imaging performance lagged behind that of CCD. Availability of near or sub-micron CMOS technology, maturity of CMOS processing, in conjunction with the advent new low noise active pixel sensor (APS) concepts have altered the situation in the decade of the nineties [1, 2, 3]. APS approach enables high quality CMOS imagers with performance rivaling those of CCDs, drawing video technology into the mainstream of CMOS system-on-a-chip development.

The primary advantages of CMOS APS are low-cost, low-power (100-1000x lower than CCDs), simple digital interface, random access, simplicity of operation (single CMOS compatible

power supply), high speed (>1000 frames per sec.), miniaturization (10-100x smaller) through system integration, and smartness by incorporating on-chip signal processing circuits. In this paper, we report the first single-chip, large format, high quality fully digital, fully programmable, ultra-low power electronic camera. We report design, operation, and performance of a camera-on-a-chip and a microminiature digital camera built with the CMOS imager chip.

IMAGER ARCHITECTURE & OPERATION

The schematic for the single-chip fully digital camera implemented with CMOS APS is shown in figure 1. This CMOS VLSI imager consists of a two-dimensional (2-D) pixel array, an array of column-parallel analog-to-digital converters (pitch-matched to the pixel array), a linear array of analog signal conditioning circuits, programmable reference generation circuits (DACs), and digital control circuits. A set of input registers enables the user to program the exposure time, electronic pan & zoom setting, speed & bias control. In addition to the imager state-machine, the digital control circuits include row & column decoders, and row logic elements for addressing pixels in the array.

Each cell of an APS consists of a photodetector, buffered by a low-noise multiplexed linear amplifier. Unlike CCD, an APS pixel produces a voltage output over the column bus, and can be selected using the row line. Figure 2 shows the schematic for the APS pixel. Figure 2a shows a photodiode (PD) type pixel consisting of a reverse-biased p-n junction for optical collection, and figure 2b shows a photogate (PG) type pixel where the photoelectrons are collected under a photogate (PG), and are separated from the floating diffusion (FD) by the isolation gate (TX). Both types of pixel occupy a small area of 12 μm x 12 μm, enabling high resolution imaging. PG pixel operation has been presented elsewhere [4]. RST, PG and SEL are row-decoded signals, and SHS and SHR are global clocks that are common to all columns.

The pixel amplifier is a simple source follower (SF) comprised of transistors Min, Msel, and

Mld. Source follower is chosen as the pixel amplifier because of its excellent gain uniformity (better than 0.5 %) enabling highly uniform image reproduction. Power dissipation in the pixel array is minimized by sharing the load transistor (Mld) to the entire column bus (BUS), so that current flows only through those source followers that belong to the selected row.

The imager operates in a column-parallel fashion (i.e. a row at a time). For readout of a photodiode pixel a row is selected, and the SF output (VS) at the column bus is sampled on the capacitor (CS) by pulsing SHS. The sampled value represents the diode potential following the collection of photoelectrons in the depletion region. Following this, FD is reset by pulsing RST, and the diode potential is read out again, and sampled on CR, producing VR. Calculation of the difference between VR and VS by temporal difference sampling (TDS) provides a measurement of the strength of optical signal received. At the same time, TDS enables cancellation of pixel-to-pixel offsets due to MOSFET threshold (V_T) mismatches, as well as reduction of flicker noise.

The imager operation is different from the one presented earlier [5]. In this imager, the imaging sequence is controlled by a read and a integrate pointers shown schematically in figure 3a. A pointer is a row-decoded and latched signal, driven by an on-chip counter. Normally, the pixel array is in reset state (digital still imaging mode). To begin imaging, the integrate pointer causes the RST clock for the start row to be pulsed and held low. This allows pixels in that row to begin exposure. The integrate pointer is scanned down at the row rate (t_{row}), while the read pointer follows it at a fixed delay, the delay being equal to exposure time (t_{int}). The read pointer causes a pixel read operation (described above), following which the pixel is returned to reset state. Exposure time is easily programmed by writing appropriate values to the delay register, the minimum exposure being equal to a row time.

Each row read consists of four non-overlapping phases as shown in the sequence diagram in figure 3b. Nominally, each phase consists of a read operation for row m, and an integrate operation for row $m \pm k$. At the end of pixel read from row m, sampled-and-held signals from

column capacitors (CS & CR) are fed to respective column ADCs. The ADC phase is followed by sequential readout of digital outputs for row m, completing the row operation. For analog readout, potentials on CS and CR are directly scanned out using in-column multiplexed source followers. The total frame time (t_{fr}) is given by:

$$t_{fr} = N \cdot t_{row} = N \cdot \left(t_{pix} + t_{adc} + N \cdot t_{rdo}\right);$$ where N is the imager format.

Frame time is primarily determined by the output clocking speed, with t_{pix} and t_{read} being a small fraction of t_{row}, especially if the digital imager is read out in serial mode. Overlapping of phases (pipelined operation) allows higher operating speed, only at the cost of unacceptable degradation of image quality due to digital noise. For typical values of t_{pix}=1.5 μs, t_{adc}=12 μs, and t_{rdo}=20 ns, video frame rate of 30 FPS is achieved even with a 1024x1024 sized imager.

ANALOG-TO-DIGITAL CONVERTER DESIGN

A column-parallel (CP) ADC architecture refers to an ADC array where each ADC serves an individual column-output. The parallelism in this architecture enables minimization of power without sacrificing data rates, by allowing individual ADCs to operate at much lower speeds. Only 200 kHz conversion rate is required for a 1024^2 imager operating at video rate, producing 30 Mpix/s output. Figure 4a shows the ADC conversion rate of such an imager as a function of the number of columns multiplexed to an ADC for different frame rates. For a fixed data rate (dictated by image acquisition requirements), the conversion rate rises non-linearly with the frame rate. Even at moderate frame rates of 15 FPS, conversion rate for a single ADC per chip (M=1024) architecture is 50 MHz. High speed high resolution ADC design is not only difficult, but is also very power hungry. Figure 4b shows the relative increase in ADC power as a function of different numbers of columns multiplexed. Each curve is normalized to the power required for CP-ADC implementation at that frame rate. The graph indicates that CP architecture results in lowest power implementation, with the power advantage improving at higher data rates. For

video-rate applications, CP architecture provides over 100x less power than a single ADC per-chip approach. The super-linear increase in ADC power is a reflection of the fact that, unlike digital power, analog bias current varies as the square of the operating speed [6]. One major concern with CP architecture is the layout challenges due to severe space restriction in column direction (limited to pixel pitch), although vertical dimension can extend to a few mm. The space constraint usually precludes the use of active circuits (e.g. opamps) for ADC implementation.

We have used an all-capacitor successive-approximation (ACSA) ADC algorithm that is not only compatible layout constraints of CP architecture, but also enables ultra low-power operation. The 10-bit column-ADC circuit schematic along with the bit cell circuit is shown in figure 5. In SA-ADC, a digital-to-analog converter (DAC) generates a weighted sum of binary-scaled reference voltage to successively approximate the signal to be digitized, at the end of which DAC input is the equivalent digital code. In ACSA implementation, the DAC is realized using a binary-scaled capacitive network, with each capacitor connected to a bit cell (Bi). Each Bi consists of a latch and logic (circuit schematic shown in figure 5b) circuits to provide V_{ref} or ground connection to the bottom plate of each capacitor. Initially reset (VR) and signal (VS) values are sampled with respect to ground respectively on CR and CS.

CS is implemented with N_b (equal to ADC resolution) binary-scaled capacitors (C_i). For symmetry CR is also split into binary-scaled capacitors C_{Ri}'s. If C is the largest sized capacitor used, each capacitor in the bank follows the relation $C_i = \frac{C}{2^i} \Rightarrow CS = 2 \cdot C$. Since the unit capacitors (C_i) in CS are binary-scaled, pulsing the bottom plate of a capacitor C_i to V_{ref} raises the voltage on CS by the capacitive divider ratio or equal to $V_{ref}/2^i$.

The ADC operation consists of successively adding the DAC generated binary-scaled voltages to the signal side (VS) to allow it to catch up to VR. If at the i-th cycle, voltage on CS is greater than that on CR, the i-th bit cell output is latched to a logic level "1" by pulsing latch-en,

and the voltage on CS is returned to its original value. Otherwise, the bit cell logic latches the output to logic "0". Digitization is completed after N_b repetitions, with bit cells representing digital code. TDS operation (reading the pixel differentially) is inherent to the ADC operation, leading to minimization of pixel-to-pixel offset variations resulting from V_T mismatches.

The bias and reference generation block consists of 6 multiplying DACs (MDACs) to provide the required reference voltages and bias currents. A 5-bit DAC is sufficient, since accurate reference is not needed, as long as they are low noise, and have sufficient drive to allow fast-settling and slew. The MDACs are realized by selectively summing outputs of binary-scaled current mirrors, the current source selection being carried out by appropriate MDAC input bits.

CONVERSION ACCURACY

The ADC inaccuracy is determined by relative errors in capacitor matching, and by voltage-dependent comparator offset. The latter can be made insignificant by buffering the comparator input, that prevents capacitive "kickback". The error in capacitor matching is reduced by careful layout and through the use of large sized capacitors. Minimum capacitance is limited to 20 fF, making the largest capacitance C=10 pF. To improve matching, each capacitor is implemented by concatenating a set of base capacitors. This method reduces relative errors both by randomization errors in base capacitor, and by minimization of fringe field mismatches through realization of constant area-to-perimeter ratio. Taking relative error into account, cell capacitance is: $\tilde{C}_i = \frac{C}{2^{i-1}} \cdot (1 + \varepsilon_i)$, where ε_i is the relative error, and $\varepsilon_1 = 0$. The maximum error occurs for the code transition [011...1] to [100...0]. The worst case differential non-linearity (DNL) will result when $\varepsilon_i = \varepsilon_{max}$. In that case, the integrated error is given by: $\varepsilon_{worst} = \left[\frac{C}{\sum_{1}^{N} \tilde{C}_i} - \frac{\sum_{2}^{N} \tilde{C}_i}{\sum_{1}^{N} \tilde{C}_i} \right] \cdot V_{ref} = \frac{\varepsilon_{max} \cdot V_{ref}}{2}$.

To achieve DNL < 0.5 LSB, maximum tolerable relative error (worst case) is 2^{-N_b} or ~ 0.1 % for a

10-bit implementation. This kind of capacitor matching is possible in CMOS VLSI, indicating that a 10-bit ACSA-ADC meets the worst case error criterion. At the other end, if the error is entirely randomized with an r.m.s. error $<\varepsilon_i>$ = ε, then the maximum DNL is: $\varepsilon_{best} = \left[\left\{\frac{C^2}{\sum_1^N \tilde{C}_i^2}\right\} - \left\{\frac{\sum_2^N \tilde{C}_i^2}{\sum_1^N \tilde{C}_i^2}\right\}\right]^{0.5} \cdot V_{ref} \approx \frac{3 \cdot \varepsilon^2 \cdot V_{ref}}{4}$. In that case, the tolerable relative error is only $\sqrt{\frac{1}{3 \cdot 2^{Nb-2}}}$ or as high as 3.6 %. Error analysis and simulations indicate that ACSA-ADC would be able to provide 10-bit digitization with DNL < 0.5 LSB.

Reduction of column-to-column ADC offset-mismatch is achieved by estimating the ADC offset during a dummy cycle, and subtracting it during regular operation. Offset is estimated by sampling the same voltage on both CS and CR, followed by digitization. Offset correction is ideally required only once, although is performed once per frame to provide drift immunity.

POWER DISSIPATION

Low average power is achived by temporarily turning off tail currents in circuit blocks with no activity. Although instantaneous power in the V_{ref} generating DAC, as well as the pixel amplifiers is high, current flow occurs only for a fraction(~ 1-10 %) of the row time. The digital control logic power is negligible, since except for column scan, the entire state machine runs off a slow row clock. In order to minimize ADC power, bit times (t_{bi}) are binary-scaled, since the equivalent capacitance (C_{bi}) reduces exponentially in successive cycles. Average power in different blocks of the camera-on-a-chip, operated in analog and digital modes, are given below.

$$P_{adc} = \frac{t_{adc}}{t_{row}} \cdot N \cdot V_{ref}^2 \cdot \Sigma \frac{C_{bi}}{t_{bi}} = 2 \cdot FPS \cdot N^2 \cdot (Nb-1) \cdot C \cdot V_{ref}^2; \quad C_{bi} = \frac{C \cdot (2C - C_i)}{2C}$$

$$P_{pix} = FPS \cdot N^2 \cdot t_{pix} \cdot I_{pix}; \quad P_{dac} = 2 \cdot FPS \cdot N^2 \cdot C \cdot V_{ref} \cdot V_{dd};$$

$$P_{dig\text{-}rdo} = 0.5 \cdot FPS \cdot N^2 \cdot C \cdot V_{dd}^2; \quad P_{anlg\text{-}rdo} = I_{out} \cdot V_{dd}$$

Figure 6 is a log-log plot of the total analog and digital power dissipation, along with power

dissipations in DAC, pixel, and ADC. The plot reveals that the average power dissipation is higher (more than 4 times at video rates) when the camera-on-a-chip is operated in analog than in digital mode, even after including ADC power. This is due to the fact that most of power dissipation occurs during readout, and for high data rates analog power varies as the square of the data rate (due to accurate settling requirements), while digital power rises only linearly. The non-linear dependence of I_{out} on frame rate causes the analog imager to be requiring more power for operation compared to a digital imager operating under same conditions. Measured power dissipation from the imager chip closely follows the relationship presented above, as shown by the open dots in figure 6, deviating from the theoretical curve only at very low frame rates. This deviation is caused by insufficient DAC resolution, preventing generation of small currents.

IMAGER NOISE

ADC noise consists mostly of switching noise associated with C_{bi}. Summing noise contributions from each switched bit cell in quadrature, the total voltage noise is estimated to be: $\frac{\pi kT}{4C}$, which is negligible (~ 15 μV), for C ~ 10 pF. For a well-designed system, the pixel noise dominates the overall noise performance of the imager chip. The pixel noise can consist of two sources: first, the reset noise due to periodic reset of the floating diffusion (FD in figure 2), and secondly, the source follower noise sampled on CS and CR. The input-referred reset noise in electrons is the familiar $\left[\frac{kT \cdot C_{FD}}{\kappa q}\right]^{0.5}$, where C_{FD} is the floating diffusion capacitance, and q is the electronic charge, κ (~ 4-5) is the noise reduction factor due to reseting a node with a transistor biased in weak-inversion, and kT/q is the thermal voltage. The reset noise is eliminated in PG-pixel by correlated double sampling (CDS), implemented by sampling VR before VS. Unlike with a PD-pixel, CDS operation is possible with a PG-pixel, since its integration node and sense node are physically separate (figure 2). The input referred noise expression due to SF is presented below, with the first term is due to white noise, and the second due to flicker noise [7].

$$\left\langle N_{in}^2 \right\rangle = \frac{2kT}{q^2}\cdot\left[\frac{(1+\beta)}{A}\cdot\left(\frac{C_{FD}+C_{gsin}}{C_{SH}}\right)\cdot C_{eq}\right] + 2\cdot\frac{(C_{FD}+C_{gsin})^2}{q^2}\cdot S_{vi}(1)\cdot(1+\beta^2)\cdot\ln(1+10\pi^2);\quad C_{eq} = C_{FD} + (1-A)C_{gsin}$$

where β is the transconductance ratio between Mld and Min, C_{SH} is the sampling capacitance (CS or CR), C_{gsin} is the gate-to-source capacitance of M_{in}, A is the SF gain, and S_{vi} (1) is the flicker noise power spectral density at 1 Hz White noise is dominant, since TDS operation suppresses flicker noise by restricting noise bandwidth through temporal correlation.

C_{eq} is the equivalent capacitance at the FD node, and is equal to the inverse of the charge-to-voltage conversion gain. For the same output noise, a smaller value of C_{eq} reduces input referred noise. Overall reduction of noise is achieved by increasing the sampling capacitor C_{SH}, since it is a sampled data system governed by kT/C type noise behavior. The equation also captures the excess noise due to capacitive feedback through C_{gsin}, increasing the noise effective bandwidth. Figure 7 plots the noise in PG- and PD-pixels as a function of S/H capacitance and C_{FD}. The plot shows that for C_{SH} ~ 3 pF, extremely low noise (< 5 e^-) is possible with PG-pixel, whereas the noise in a PD-pixel (~ 27 e^-), though higher, is quite acceptable for most commercial and space applications. Measured noise (open dots in figure 7) closely follows the theoretical curve. Accuracy of noise measurement from PG-pixel is impeded by relatively higher system noise, causing a variation from the theoretical predictions. Since the noise behavior is governed by the pixel amplifier noise, and since pixel power dissipation is negligible, APS imager noise, unlike CCD, is relatively independent of data rate, enabling high-speed, low-noise applications.

IMAGER PERFORMANCE

Figure 8 shows the die photograph of the 512x512 digital camera-on-a-chip. The chip occupies an area of 10 mm x 12 mm, with the imager area being 6 mm x 6 mm, and the ADC area is 6 mm x 5 mm. The pixel pitch is 12 μm, with drawn optical fill-factor of 38%. CP-ADC array takes up a large amount of area since accurate capacitance matching is desired. The chip provides digital video data (8 Mpix/s) at an ultra-low power of ~ 15 mW, while producing high quality

image. Figure 9 is a reproduction of the image of a portion of a dollar bill captured with the digital camera-on-a-chip. Unlike other CMOS imagers, no image artifacts such as column-to-column fixed pattern noise (FPN) are visible, and the pixel response non-uniformity is less than 1 % (typical of a camcorder-type imager). The column-to-column FPN, measured from a flat field, is less than 1.5 LSB across the array under low-light levels, and can be further reduced by software corrections. The imager shows good optical response with excellent linearity, shown in figure 10. Better than 9 bit linearity is achieved over 90 % of the full-well, the residual non-linearity being contributed by source-followers and non-linear diode capacitance. The peak quantum efficiency at 550 nm is 42 %, indicating good optical collection. Figure 11 shows the differential non-linearity plot of the ADC. As expected, most of the non-linearity shows up at major transition codes, but is less than $\pm$ 0.5 LSB over the entire input range. The CMOS imager performance, summarized in table 1, is commensurate with high quality imaging requirements.

Figure 12 shows the photograph of the first-ever fully digital micro-miniature camera implemented with CMOS digital camera-on-a-chip technology. The total volume of the camera including optics is 3 cm^3, and is comprised mostly of optics. The optics size can be further reduced, allowing "wrist-watch" type cameras. The micro-miniature programmable camera is made possible by realization of a single-chip digital imager that interfaces with only 4 wires (Vdd, Gnd, clock, and serial Din/Dout). A similar miniature camera will be used for autonomous rover navigation in the next MARS mission.

CONCLUSIONS

In summary, we have presented for the first time design and performance of a large format, ultra-low power, single-chip digital imager capable of reproducing high quality images, and implemented in CMOS technology. We have presented the design considerations and trade-offs for reproduction of high quality images in CMOS technology, while minimizing power and increasing data rates. The 512x512 digital imager requires a single power supply (3.3 V),

consumes only 15 mW of power at video rates, is capable of operation with read noise < 30 e^- in photodiode mode, & < 5 e^- in photogate mode. Demonstration of high quality imaging in CMOS technology has opened the door for future developments of smart and integrated imaging system-on-a-chip.

REFERENCES

1. S. Mendis, B. Pain, S. Kemeny, R. Gee, Q. Kim, E. Fossum, *CMOS active pixel image sensors for highly integrated imaging systems*, IEEE J. of Solid-state Circuits, vol. 32 (2), pp. 187-198, 1997.

2. S. Smith, J. Hurwitz, M. Torrie, D. Baxter, A. Holmes, M. Panaghiston, R. Henderson, A. Murray, S. Anderson, P. Denyer, *A single chip 306x244-pixel CMOS NTSC video camera*, Tech. Digest, IEEE ISSSC, vol. 41 (FA 11.2), pp. 170-171, 1998.

3. H. Ihara, H. Yamashita, I. Inoue, T. Yamaguschi, N. Nakamura, H. Nozaki, *A 3.7x37 μm^2 square pixel CMOS imager sensor for digital still camera application*, Tech. Digest, IEEE ISSSC, vol. 41 (FA 11.8), pp. 182-183, 1998.

4. S. Mendis, S. Kemeny, B. Pain, E. Fossum, *Progress in CMOS active pixel sensors*, in Charge-coupled Devices and Solid-state Optical Sensors IV, Proc. SPIE, vol. 2172, pp. 19-29, 1994.

5. R. H. Nixon, B. Pain, S.E. Kemeny, C. O . Staller, and E.R. Fossum, *256x256 CMOS active pixel sensor camera-on-a-chip*, IEEE Journal of solid-state circuits, vol. 31 (12), pp. 2046-2051, 1996.

6. B. Pain and E.R. Fossum, *Approaches and analysis for on-focal-plane analog-to-digital conversion*, in Infrared Readout Electronics II, Proc. SPIE, vol. 2226, pp. 208-218, 1994.

7. B. Pain, O.Yadid-Pecht, B. Mansoorian, E. Fossum, *Optimization of active pixel sensor noise and responsivity for scientific applications*, in Solid State Sensor Arrays: Development and Applications, Proc. SPIE vol. 3019, pp. 125-136, 1997.

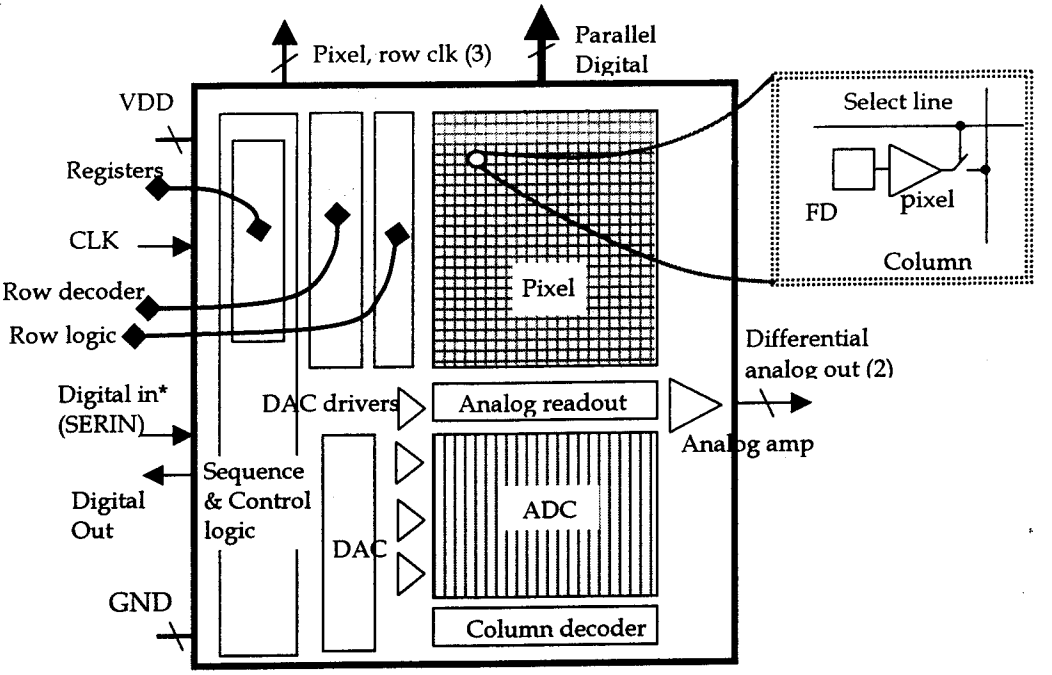

Figure 1: Schematic of the digital camera-on-a-chip

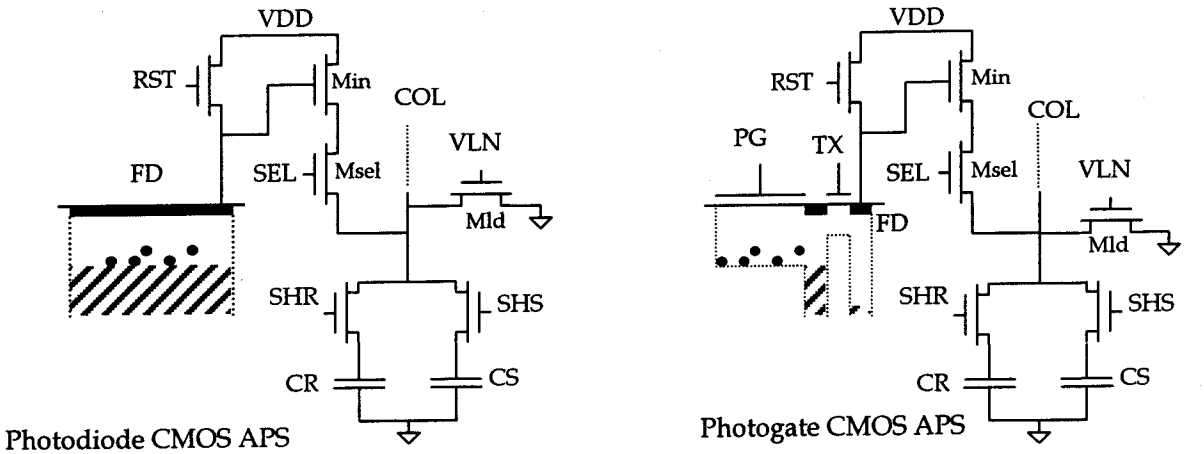

Figure 2: CMOS active pixel sensor unit cell. 2a: photodiode pixel; 2b: photogate pixel

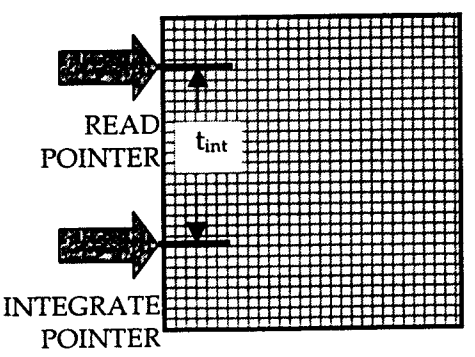

Figure 3a: Chip architecture with integrate & read pointers

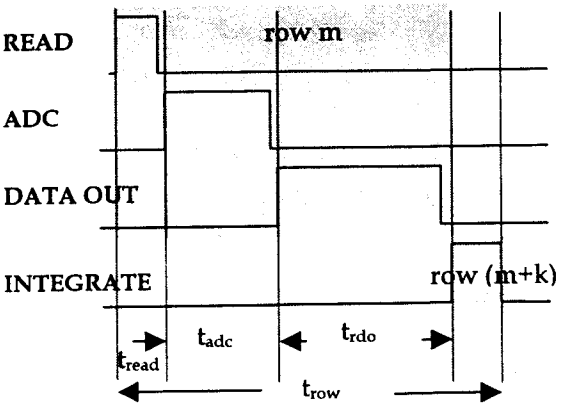

Figure 3b: Imager sequence phase diagram

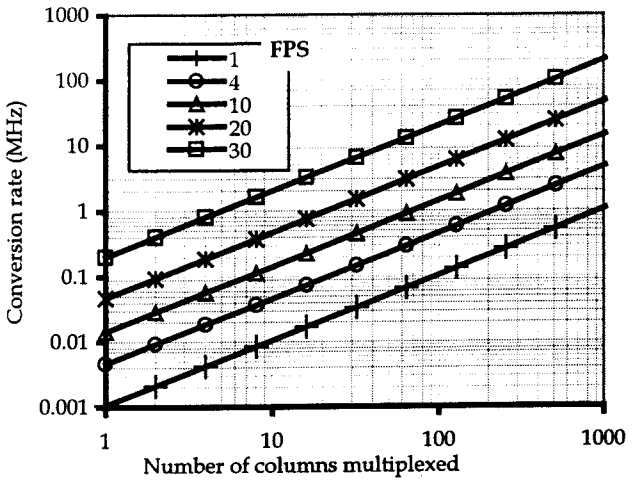

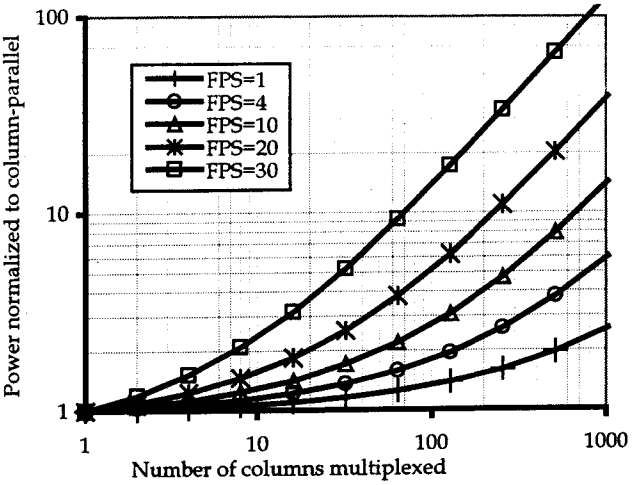

Figure 4a: ADC conversion rate vs. number of columns multiplexed (1024x1024 format);

Figure 4b: ADC power normalized to power for column-parallel approach at different frame rates

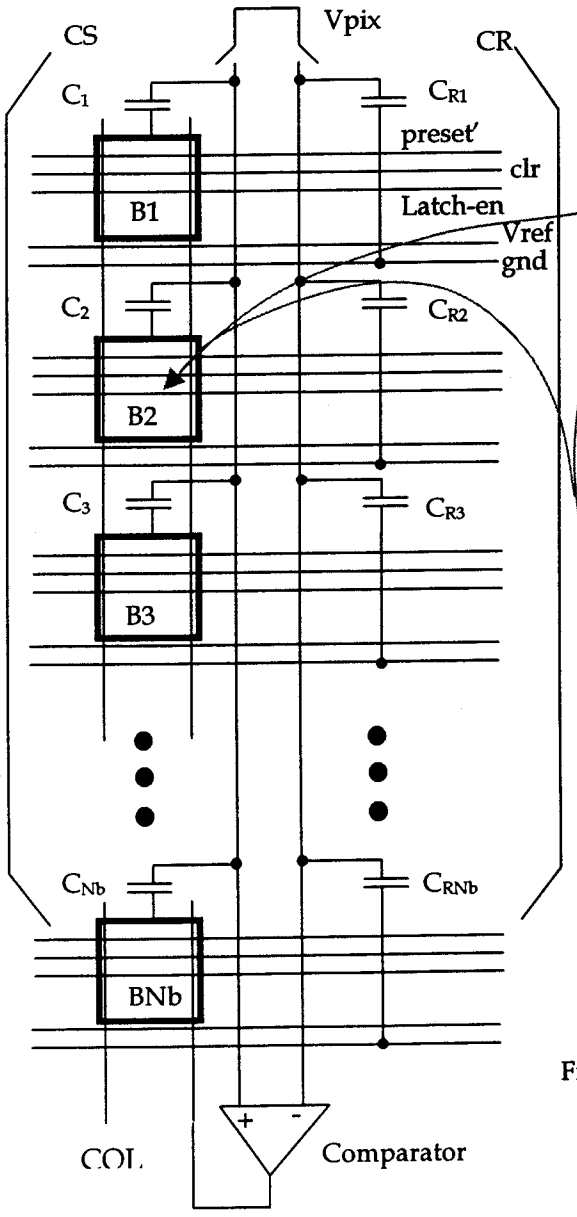

Figure 5a: Schematic of the ACSA-ADC

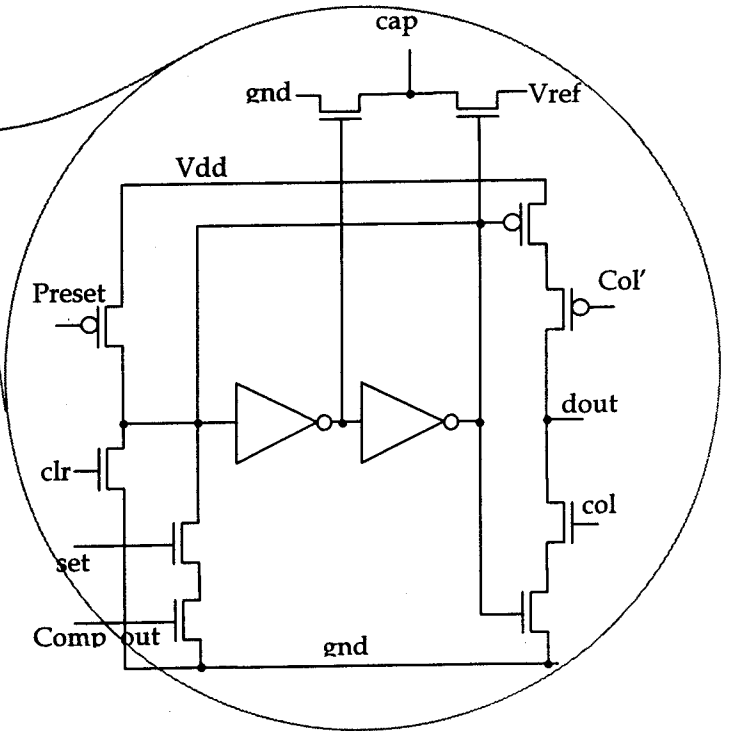

Figure 5b: Schematic of a bit cell

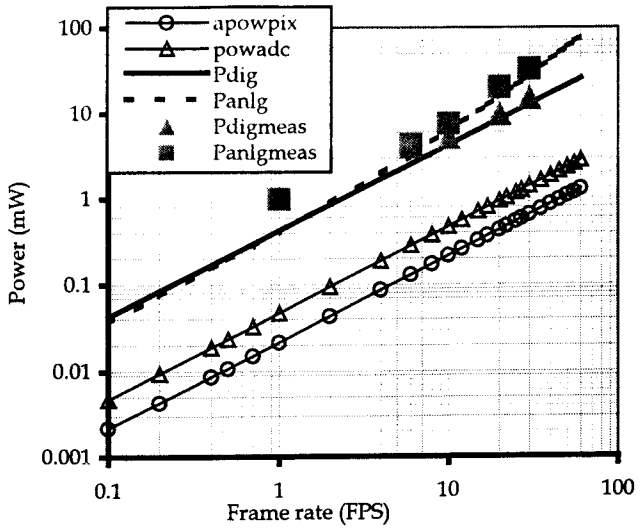

Figure 6: Power dissipation as function of frame rate

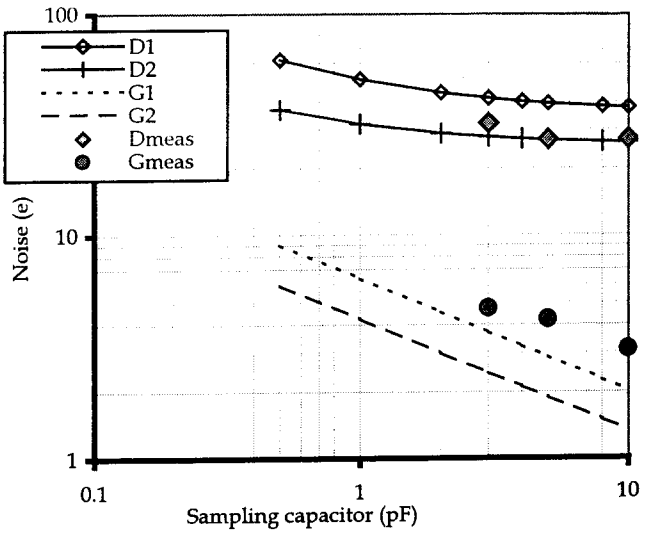

Figure 7: Input Noise as function of C_{SH}

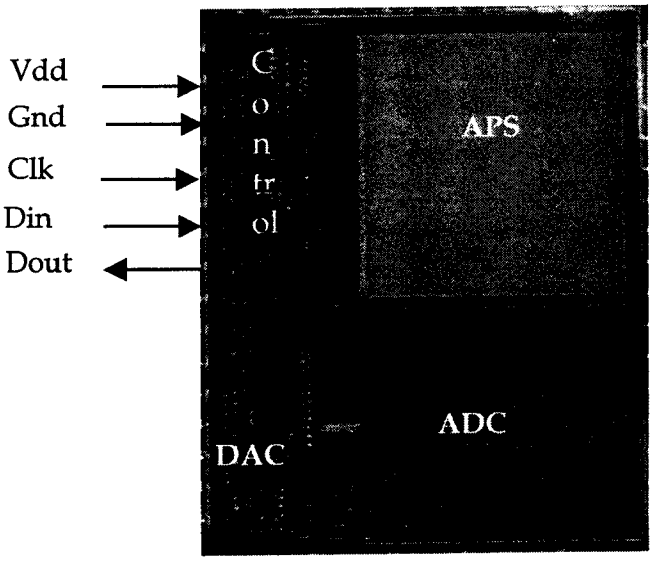

Figure 8: Digital camera-on-a-chip photo

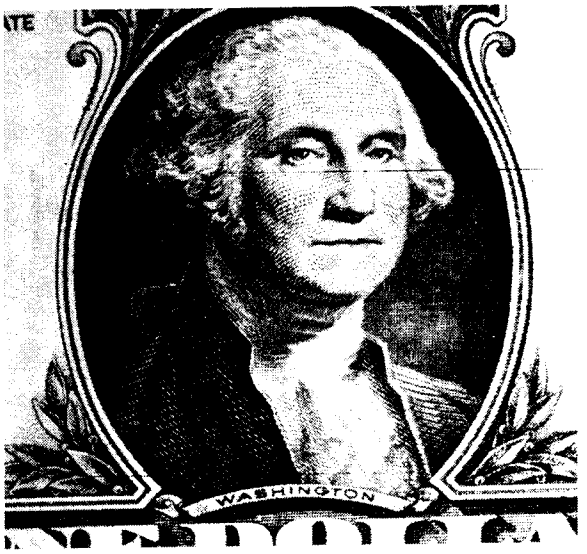

Figure 9: Reproduced image of "George"

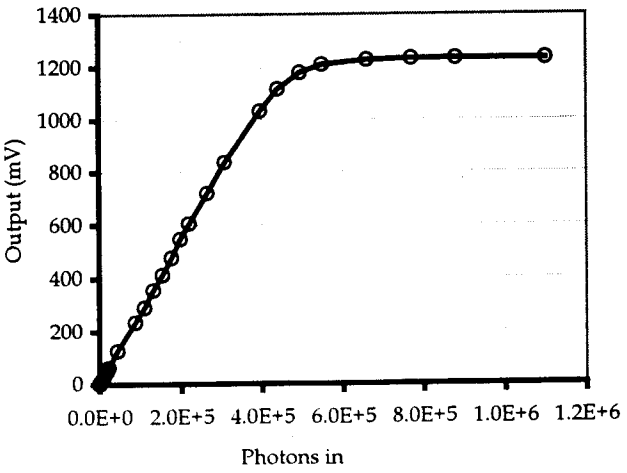

Figure 10: Photon transfer curve

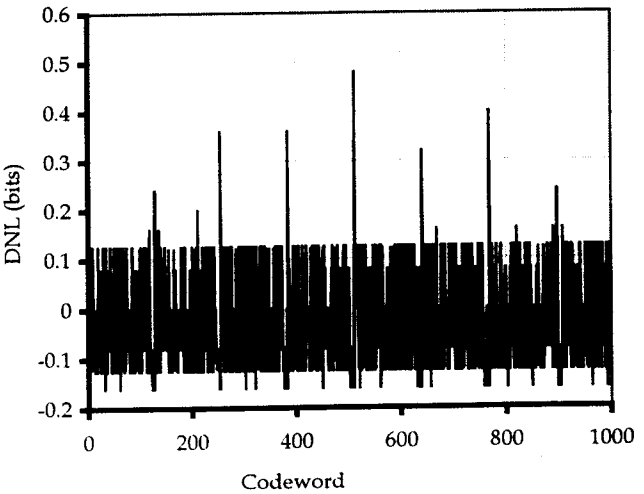

Figure 11: Differential non-linearity of ADC array

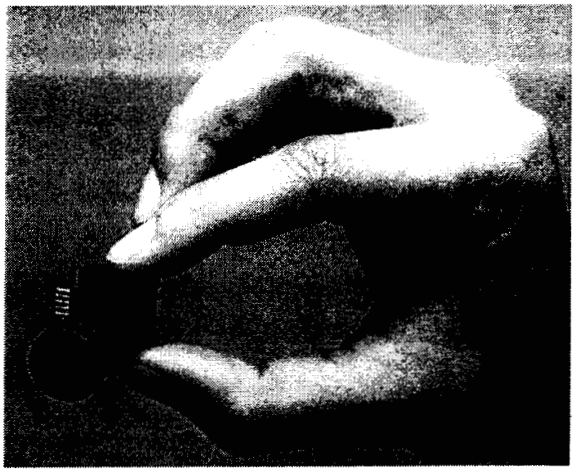

Figure 12: Picture of micro-miniature 4-wire digital camera

Table 1: APS camera-on-a-chip performance summary

Imager characteristics	Values	Comments
Technology	CMOS, 0.6 μm	
Chip size	10 mm x 12 mm	
Outputs	Analog & digital	
Format	512 x 512	
Pixel size	12 μm x 12 μm	
Optical fill factor	37 % for PD	
Responsivity	3 μV/photon	@ 550 nm
Linearity	> 99.85 %	
Quantum efficiency	42 % for PD	For the total pixel
Dark current	500 pA/cm² for PD	
Fixed pattern noise	< 0.1 % saturation	
Noise	26 e^- for PD 5 e^- for PG	
Full well	350,000 e^- for PD 48,000 e^- for PG	
Analog dynamic range	82.6 dB for PD 79.6 db for PG	
ADC resolution	10 bits	
ADC DNL max.	0.46 LSB	
Power	15 mW	@ 10 Mbits/sec.
Power supply	3.3V	Single power supply
Timing and control	on-chip	
Bias generation	on-chip	
Max. frame rate	35 FPS	